arXiv:1603.00847v4 [cs.CG] 4 Feb 2017

Shortest Paths and Convex Hulls in 2D Complexes with Non-Positive Curvature

Anna Lubiw* Daniela Maftuleac* Megan Owen†

Abstract

Globally non-positively curved, or CAT(0), polyhedral complexes arise in a number of applications, including evolutionary biology and robotics. These spaces have unique shortest paths and are composed of Euclidean polyhedra, yet many algorithms and properties of shortest paths and convex hulls in Euclidean space fail to transfer over. We give an algorithm for computing convex hulls using linear programming for 2-dimensional CAT(0) polyhedral complexes with a single vertex. We explore using shortest path maps to answer single-source shortest path queries in 2-dimensional CAT(0) polyhedral complexes, and we improve and unify solutions for 2-manifold and rectangular cases.

1 Introduction

Convex hulls and shortest paths—and algorithms to find them—are very well understood in Euclidean spaces, but less so in non-Euclidean spaces. We consider these two problems in finite *polyhedral complexes* which are formed by joining a finite number of d-dimensional convex polyhedra along lower dimensional faces. We will primarily be concerned with the 2D case of triangles or rectangles joined at edges.

We will restrict our attention to polyhedral complexes that are globally non-positively curved, or CAT(0). Introduced by Gromov in 1987 [25], CAT(0) metric spaces (or spaces of global non-positive curvature) constitute a far-reaching common generalization of Euclidean spaces, hyperbolic spaces and simple polygons. The initials "CAT" stand for Cartan, Alexandrov, and Toponogov, three researchers who made substantial contributions to the theory of comparison geometry. In a CAT(0) space, in contrast to a space of positive curvature, there is a unique geodesic (locally shortest path) between any two points and this property characterizes CAT(0) complexes.

*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada, alubiw@uwaterloo.ca,daniela.maftuleac@uwaterloo.ca; † Department of Mathematics and Computer Science, Lehman College, City University of New York, United States, megan.owen@lehman.cuny.edu

The impact of CAT(0) geometry on mathematics is significant especially in the field of geometric group theory where the particular case of CAT(0) polyhedral complexes formed by cubes—the so-called "CAT(0) cube complexes"—are particularly relevant [10, 25, 28]. Most of the work on CAT(0) metric spaces so far has been mathematical. Algorithmic aspects remain relatively unexplored apart from a few results for some particular CAT(0) spaces [16, 17, 21].

This paper is about algorithms for finite CAT(0) polyhedral complexes, which we will call "CAT(0) complexes" from now on. We are primarily interested in the algorithmic properties of CAT(0) complexes because of their applications, particularly to computational evolutionary biology. The (moduli) space of all phylogenetic (evolutionary) trees with n leaves can be modelled as a CAT(0) cube complex with a single vertex [7], and being able to compute convex hulls in this space would give a method for computing confidence intervals for sets of trees (see Section 2.3 for more details). A second application of CAT(0) cube complexes is to reconfigurable systems [24], a large family of systems which change according to some local rules, e.g. robotic motion planning, the motion of non-colliding particles in a graph, and phylogenetic tree mutation, etc. In many reconfigurable systems, the parameter space of all possible positions of the system can be seen as a CAT(0) cube complex [24].

Main Results. In this paper we study the shortest path problem and the convex hull problem in 2D CAT(0) complexes formed by triangles or rectangles. For any set of points P we define the *convex hull* to be the minimal set containing P that is closed under taking the shortest path between any two points in the set. We show that convex hulls in 2D CAT(0) complexes fail to satisfy some of the properties we take for granted in Euclidean spaces. Our main result is an algorithm to find the convex hull of a finite set of points in a 2D CAT(0) complex with a single vertex. In general, for any CAT(0) complex, the convex hull of a set of points is the union of a convex set in each cell of the complex. For the case of 2D CAT(0) complexes, these convex sets are polygons. For the special case when there is a single vertex, we show how to find these polygons using linear programming. Our algorithm runs in polynomial time (in bit complexity) for a cube complex. For more general inputs we must use the real-RAM model of computation, and the bottleneck in our running time is the time required for linear programming in an algebraic model, which is not known to be polynomially bounded, but is considered efficient via the simplex method.

In the single-source shortest path problem, we are given a 2D CAT(0) complex of n triangles and a source point s, and we wish to preprocess the complex in order to find the shortest path from s to any query point t quickly. We explore the shortest path map, which divides the space into regions where shortest paths from s are combinatorially the same (i.e. traverse the same sequence of edges and faces). We show that the shortest path map may have exponential size. An alternative, the "last step shortest path map," has linear size and can be used to find shortest paths from s in time proportional to the number of faces traversed by the path. We show how to construct the last step shortest path map in $O(n^2)$ preprocessing time and space for special cases. This generalizes two previous results: an algorithm by Chepoi and Maftuleac [18] for the case of 2D CAT(0) rectangular complexes; and an algorithm by Maftuleac [39] for the case of a 2D CAT(0) complex that is a *topological*

2-manifold with boundary (i.e. every edge is incident to at most two faces). In the latter case we improve the time bounds.

The rest of the paper is organized as follows. Section 2 contains further background on the problem, including existing algorithmic results for CAT(0) polyhedral complexes and applications to phylogenetics. Section 3 reviews the relevant mathematics and tree space notation. Section 4 gives our results for convex hulls, and Section 5 gives our results for shortest paths in 2D CAT(0) polyhedral complexes. Finally we give our conclusions in Section 6.

2 Background

In this section we describe background work on shortest path and convex hull algorithms, and discuss the application of our work to phylogenetic trees.

One of the most basic CAT(0) spaces is any simple polygon (interior plus boundary) in the plane. This can be viewed as a 2D CAT(0) complex once the polygon is triangulated. The fact that geodesic paths are unique is at the heart of efficient algorithms for shortest paths and related problems. On the other hand, generalizing a polygon to a polygonal domain (a polygon with holes) or a polyhedral terrain yields spaces that are not CAT(0), since geodesic paths are no longer unique. This helps explain why shortest path and convex hull problems are more difficult in these more general settings.

2.1 Shortest Paths

The shortest path problem is a fundamental algorithmic problem with many applications, both in discrete settings like graphs and networks (see, e.g., [1]) as well as in geometric settings like polygons, polyhedral surfaces, or 3-dimensional space with obstacles (see, e.g., Mitchell [41]).

All variants of the shortest path problem can be solved efficiently for a polygon once it is triangulated, and triangulation can be done in linear time with Chazelle's algorithm [13]. The shortest path (the unique geodesic) between two given points can be found in linear time [36]. For query versions, linear space and linear preprocessing time allow us to answer single-source queries [26,29], and more general all-pairs queries [27], where answering a query means returning the distance in logarithmic time, and the actual path in time proportional to its number of edges.

By contrast, in a polygonal domain, where geodesic paths are no longer unique, the best shortest path algorithm uses a continuous-Dijkstra approach in which paths are explored by order of distance. For a polygonal domain of n vertices, this method takes $O(n \log n)$ [preprocessing] time [30] (see the survey by Mitchell [41]). For a polyhedral terrain the continuous-Dijkstra approach gives $O(n^2 \log n)$ time [42], and the best-known run-time of $O(n^2)$ is achieved by Chen and Han's algorithm [15] that uses a breadth-first-search approach.

There are no shortest path algorithms for the general setting of CAT(0) polyhedral complexes, although there are some for certain specializations. For 2D CAT(0) complexes that

are 2-manifolds, Maftuleac [39] gave a Dijkstra-like algorithm for the single-source shortest path query problem, with preprocessing time $O(n^2 \log n)$, space $O(n^2)$ and query time proportional to the size of the output path.

There are also some partial results on finding shortest paths when we restrict the CAT(0) polyhedral complex to be composed of cubes or rectangles. The space of phylogenetic trees mentioned in the introduction is a CAT(0) cube complex. For these "tree spaces," Owen and Provan [46] gave an algorithm to compute shortest paths (geodesics) with a running time of $O(m^4)$, where m is the dimension of the maximal cubes. The algorithm is much faster in practice for realistic trees. The result was extended to a polynomial time algorithm for computing geodesics in any *orthant space* [40], where an *orthant space* is a CAT(0) cube complex with a single vertex. The algorithm for tree space was also adapted by Ardila et al. [2] to compute the shortest path between two points in a general CAT(0) cube complex, however this algorithm is likely not polynomial time. Finally, Chepoi and Maftuleac [18] used different methods to give a polynomial time algorithm for all-pair shortest path queries in any 2D CAT(0) rectangular complex, with preprocessing time $O(n^2)$, space $O(n^2)$, and query time proportional to the size of the output path.

2.2 Convex Hulls

The problem of computing the convex hull of a set of points is fundamental to geometric computing, especially because of the connection to Voronoi diagrams and Delaunay triangulations [37].

The convex hull of a set of points in the plane can be found in provably optimal time $O(n \log h)$ where n is the number of points and h is the number of points on the convex hull. The first such algorithm was developed by Kirkpatrick and Seidel [35] and a simpler algorithm was given by Chan [12]. An optimal algorithm for computing convex hulls in higher dimension d in time $O(n^{\lfloor \frac{d}{2} \rfloor})$ was given by Chazelle [14]. See the survey by Seidel [48]

A simple polygon, triangulated by chords, is the most basic example of a 2D CAT(0) complex. In this setting, the convex hull of a set of points P (i.e. the smallest set containing P and closed under taking geodesics) is referred to as the *relative* (or *geodesic*) convex hull. Toussaint gave an $O(n \log n)$ time algorithm to compute the relative convex hull of a set of points in a simple polygon [49], and studied properties of such convex hulls [50]. Ishaque and Tóth [32] considered the case of line segments that separate the plane into simply connected regions (thus forming a CAT(0) space) and gave an semi-dynamic algorithm to maintain the convex hull of a set of points as line segments are added and points are deleted.

Moving beyond polygons to polygonal domains or terrains, geodesic paths are no longer unique, so there is no single natural definition of convex hull (one could take the closure under geodesic paths, or the closure under shortest paths). We are unaware of algorithmic work on these variants.

However, in a polyhedral surface with unique geodesics the convex hull is well defined, and Maftuleac [39] gave an algorithm to compute the convex hull of a set of points in $O(n^2 \log n)$ time, where n is the number of vertices in the complex plus the number of points in the set.

In all the above cases the boundary of the convex hull is composed of segments of shortest paths between the given points, which—as we shall see in Section 4.1—is not true in our setting of 2D CAT(0) complexes.

Beyond polyhedral complexes, convex hulls become much more complicated. Indeed it is still an open question if the convex hull of 3 points on a general Riemannian manifold of dimension 3 or higher is closed [6, Note 6.1.3.1]. Bowditch [9] and Borbély [8] give some results for convex hulls on manifolds of pinched negative curvature, but our complexes need not be manifolds. In the space of positive definite matrices, which is a CAT(0) Riemannian manifold, Fletcher et al. [23] give an algorithm to compute generalized convex hulls using horoballs, which are generalized half-spaces. Lin et al. [38] look at convex hulls of three points in a CAT(0) cubical complex that generalizes the space of phylogenetic trees. They prove that there are such spaces where the top-dimensional cells have dimension $2d$, and there exist 3 points in the space such that their convex hull contains a d-dimensional simplex. Bridson and Haefliger [10, Proposition II.2.9] give conditions for when the convex hull of three points in a CAT(0) space is "flat", or 2-dimensional. Finally, for a survey of convexity results in complete CAT(0) (aka Hadamard) spaces, of which the space of phylogenetic trees is one, see [4].

2.3 Application to phylogenetic trees

While this paper will look at arbitrary 2D complexes with non-positive curvature, our work is motivated by a particular complex with non-positive curvature, namely the space of phylogenetic trees introduced by Billera, Holmes, and Vogtmann [7], called the *BHV tree space*, and described in more detail in Section 3.1. Phylogenetic trees are ubiquitous in biology, and each one depicts a possible evolutionary history of a set of organisms, represented as the tree's leaves. Once we fix a set of leaves, the BHV tree space is a complex of Euclidean orthants (the higher dimensional version of quadrants and octants), in which each point in the space represents a different phylogenetic tree on exactly that set of leaves.

One area of active phylogenetics research is how to statistically analyze sets of phylogenetic trees on the same, or roughly the same, set of species. Such sets can arise in various ways: from sampling a known distribution of trees, such as that generated by the Yule process [52]; from tree inference programs, such as the posterior distribution returned by performing Bayesian inference [47] or the bootstrap trees from conducting a maximum likelihood search [22]; or from improvements in genetic sequencing technologies that lead to large sets of gene trees, each of which represents the evolutionary history of a single gene, as opposed to the species' evolutionary history as a whole. Traditionally, most of the research in this area focused on summarizing the set of trees, although recent work has included computing variance [3, 40] and principal components [44].

It is an open question to find a good way to compute confidence regions for a set of phylogenetic trees. While work has been done on proving a central limit theorem for the BHV tree space [5], even the equivalent of a Gaussian distribution is not fully known on the BHV tree space, and thus this approach can not yet be used to compute confidence regions. In [31], Holmes proposes applying the data depth approach of Tukey [51] for Euclidean space

to the BHV tree space, namely peeling convex hulls. The convex hull is the minimum set that contains all the data points, as well as all geodesics between points in the convex hull. By peeling convex hulls, we mean to compute the convex hull for the data set, and then remove all data points that lie on the convex hull. This can then be repeated. To get the 95% confidence region, for example, one would remove successive convex hulls until only 95% of the original data points remain.

If we keep peeling convex hulls until all remaining points lie on the boundary of the convex hull, then we can take their Fréchet mean [3, 40] to get an analog of the univariate median of Tukey [51]. This could also be a useful one-dimensional summary statistic for a set of trees. Many of the most-used tree summary statistics have a tendency to yield a degenerate or non-binary tree, which is a tree in which some of the ancestor relationships are undefined. This is considered a problem by biologists, but such a univariate median tree found by peeling convex hulls would likely be binary if all trees in the data set are.

Currently, these methods cannot be used, because it is not known how to compute convex hulls in BHV tree space. We show several examples of how Euclidean intuition and properties for convex hulls do not carry over to convex hulls in the BHV tree space. Our algorithm to find convex hulls applies to the space of trees with five leaves which is described in more detail in Section 3.1.

3 Preliminaries

A metric space (X, d) is *geodesic* if every two points $x, y \in X$ are connected by a locally shortest or *geodesic* path. A geodesic metric space X is *CAT(0)* if its triangles satisfy the following CAT(0) inequality. For any triangle ABC in X with geodesic segments for its sides, construct a *comparison triangle* $A'B'C'$ in the Euclidean plane such $d(A, B) = |A'B'|$, $d(A, C) = |A'C'|$, and $d(B, C) = |B'C'|$. Let Y be a point on the geodesic between B and C, and let Y' be a *comparison point* on the line between B' and C' such that $d(B, Y) = |B'Y'|$. Then triangle ABC satisfies the CAT(0) inequality if $d(Y, A) \leq |Y'A'|$ for any Y. Intuitively, this corresponds to all triangles in X being at least as skinny as the corresponding triangle in Euclidean space (Figure 1).

A *polyhedral complex* is a set of convex polyhedra ("cells") glued together by isometries along their faces. In this paper we only consider finite polyhedral complexes. When all of the cells are cubes, then this is called a *cube*, or *cubical*, *complex*. The length of a path between two points in a polyhedral complex is the sum of the Euclidean lengths of the pieces of the path in each cell of the complex. The *distance* between two points is defined as the length of the shortest path between them.

We will consider polyhedral complexes that are CAT(0). The cells are 2D (planar) convex polygons. These can always be triangulated, so the general setting is when the cells are triangles. We call this a *2D CAT(0) complex*. Sometimes we will consider a complex in which all the cells are rectangles, either bounded or unbounded. We call this a *2D CAT(0) rectangular complex*. In general we allow the space to have boundary (i.e., an edge that is incident to only one cell).

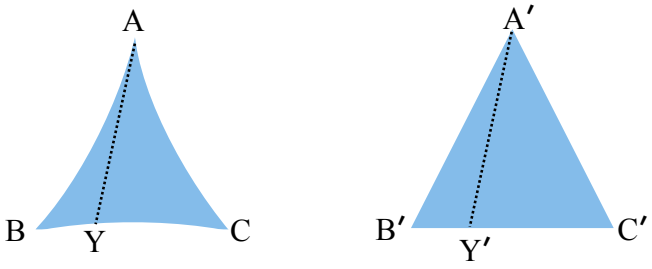

Figure 1: The triangle on the left represents a triangle in a CAT(0) space, with its corresponding comparison triangle in Euclidean space on the right.

For any vertex v of a 2D complex, we define the *link graph*, G_v as follows. The vertices of G_v correspond to the edges incident to v in the complex. The edges of G_v correspond to the cells incident to v in the complex: if r and s are edges of cell C with r and s incident to v, then we add an edge between vertices r and s in G_v with weight equal to the angle between r and s in C. Every point $p \neq v$ in cell C can be mapped to a point on edge (r, s) in G_v: if the angle between vr and vp in C is α then p corresponds to the point along edge (r, s) that is distance α from r. We use the same notation for a point of C and the corresponding point in the link graph when the meaning is clear from the context.

When we have a 2D polyhedral complex, there is also an alternative condition for determining whether it is CAT(0).

Theorem 1 ([10, Theorem II.5.4 and Lemma II.5.6])**.** *A 2D polyhedral complex $\mathcal{K}$ is CAT(0) if and only if it is simply connected and for every vertex $v \in \mathcal{K}$, every cycle in the link graph G_v has length at least 2π.*

Some of our results are only for the case where the CAT(0) complex $\mathcal{K}$ has a single vertex O, which we call the *origin*. We will call such a complex a *single-vertex complex*. In a 2D single-vertex complex every cell is a *cone* formed by two edges incident to O with angle at most π between them. There is a single link graph $G = G_O$. Every point of the complex except O corresponds to a point of G, and every point of G corresponds to a ray of points in the complex. We can decide whether the geodesic path between two points goes through O by looking at the shortest path between the corresponding points in the link graph, as follows. Consider two points a and b in $\mathcal{K}$, distinct from O, and consider the corresponding points a and b in G (context will distinguish a in $\mathcal{K}$ from a in G). Let $\sigma(a, b)$ be the (unique) geodesic path between a and b in the space $\mathcal{K}$. Let $\sigma_G(a, b)$ be a shortest path between a and b in G. Let $|\sigma|$ indicate the length of path σ.

Property 1. *Exactly one of the following two things holds:*

- $|\sigma_G(a,b)| \geq \pi$ *and* $\sigma(a,b)$ *goes through* O*,*
- $|\sigma_G(a,b)| < \pi$ *and* $\sigma(a,b)$ *maps to* $\sigma_G(a,b)$ *and does not go through* O*.*

For an example, see Figure 8. Compare the pair p_1, c, where $|\sigma_G(p_1,c)| = 160°$ and $\sigma(p_1,c)$ does not go through the origin, with the pair p_1, b, where $|\sigma_G(p_1,b)| = 200°$ and $\sigma(p_1,b)$ goes through the origin.

Proof. If $\sigma(a,b)$ does not go through O, then $\sigma(a,b)$ travels through some cells, and, by placing these one after another in the plane, $\sigma(a,b)$ forms a straight line segment through the cells, which creates a triangle together with point O. The angle of this triangle at O is $|\sigma_G(a,b)|$ which is therefore less than π.

Conversely, if $|\sigma_G(a,b)| < \pi$, then the path $\sigma_G(a,b)$ follows segments of the link graph which correspond to cells of K, and when we place these cells one after another in the plane, the angle between segments Oa and Ob in the plane is $|\sigma_G(a,b)|$. Thus the straight line segment from a to b remains in the cells, and forms a geodesic path from a to b that does not go through O, and that maps to $\sigma_G(a,b)$. □

We will make use of this property in the section on convex hulls.

3.1 BHV Tree Space

As explained in Section 2.3 the work on computing convex hulls was motivated by the BHV tree space for trees with 5 leaves, which is a 2D CAT(0) complex. We will now describe this space, which is denoted $\mathcal{T}_5$, and which contains all unrooted leaf-labelled, edge-weighted phylogenetic trees with 5 leaves (equivalently all such *rooted* trees with 4 leaves). For a description of the BHV tree spaces for trees with more than 5 leaves, see [7]. This section is not necessary for understanding the rest of the paper.

A *phylogenetic tree* is a tree in which each interior vertex has degree ≥ 3 and there is a one-to-one labelling between the leaves (degree 1 vertices) and some set of labels $\mathcal{L}$. For this paper, we assume $\mathcal{L} = \{1,2,3,4,5\}$. Also, the trees have a positive weight or length on each *interior edge*, which is an edge whose vertices have degree ≥ 3 (that is, are not leaves). If a phylogenetic tree contains only vertices of degree 1 and 3, then it is called *binary*.

A *split* is a partition of the leaf-set $\mathcal{L}$ into two parts $L_1 \cup L_2 = \mathcal{L}$ such that $|L_i| \geq 2$ for $i = 1, 2$. We write a split as $L_1|L_2$. Each interior edge in a phylogenetic tree corresponds to a unique split, where the two parts are the sets of leaves in the two subtrees formed by removing that edge from the tree. Binary trees with five leaves contain two interior edges, and hence two splits. There are 10 possible splits, and they can be combined to form 15 different tree shapes. The *shape* of a tree is defined to be the set of interior edges or splits in that tree, and tell us which species are most closely related. (Figure 2).

We now define the space $\mathcal{T}_5$ itself, which consists of exactly one Euclidean quadrant for each of the 15 possible binary tree shapes. For each quadrant, the two axes are labelled by the two splits in the tree shape. A point in the quadrant corresponds to the tree with that

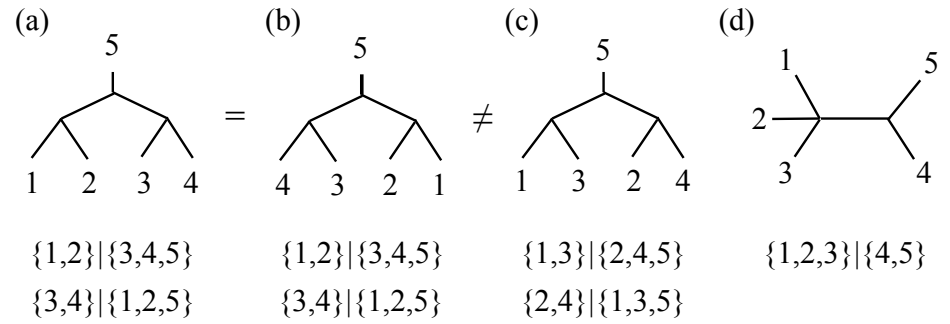

Figure 2: Several tree shapes with their constituent splits listed below them. Tree (a) and tree (b) have the same tree shape, which does not depend on the planar embedding of the tree. Tree (c) has different splits and hence a different tree shape from trees (a) and (b). Tree (d) is a non-binary tree shape.

tree shape whose interior edges have the lengths given by the coordinates. We identify axes labelled by the same split, so that if two quadrants both contain an axis labelled by the same split, then they are glued together along that shared axis (see Figure 3).

The length of a path between two trees in $\mathcal{T}_5$ is the sum of the lengths of the restriction of that path to each quadrant in turn, where it is computed using the Euclidean metric. The *BHV distance* is the length of the shortest path, or geodesic, between the two trees (Figure 3). Billera et al. [7] proved that this tree space is a CAT(0) cube complex, which implies that there is a unique geodesic between any two trees in the tree space.

To understand how the 15 quadrants in $\mathcal{T}_5$ are connected, consider the link graph of the origin, which is shown in Figure 4 and is the Petersen graph. The Petersen graph has multiple overlapping 5-cycles, one of which corresponds to the 5 quadrants in Figure 3. Also note that each vertex in the link graph of the origin is adjacent to three edges. This corresponds to each axis lying in three quadrants in $\mathcal{T}_5$. Note that this example illustrates that the link graph of even a CAT(0) rectangular complex with a single vertex need not be planar.

4 Convex Hulls

Let P be a finite set of points in a CAT(0) complex $\mathcal{K}$. Recall from Section 1 that the convex hull of P is defined to be the minimal set containing P that is closed under taking the shortest path between any two points in the set. Let $\mathrm{CH}(P)$ denote the convex hull of P in $\mathcal{K}$. For algorithmic purposes, there are several ways to specify $\mathrm{CH}(P)$. One possibility is to specify the intersection of $\mathrm{CH}(P)$ with each cell of the complex. In the case of 2D CAT(0) complexes, each such set is a convex polygon, which can be given by its vertices. Our algorithm will do this. However, we note that there is another way to specify $\mathrm{CH}(P)$, which might be easier but would still suffice for many applications, and that is to give an algorithm to decide if a given query point of $\mathcal{K}$ is inside $\mathrm{CH}(P)$.

Convex hulls in CAT(0) complexes are something of a mystery. It is not known, for

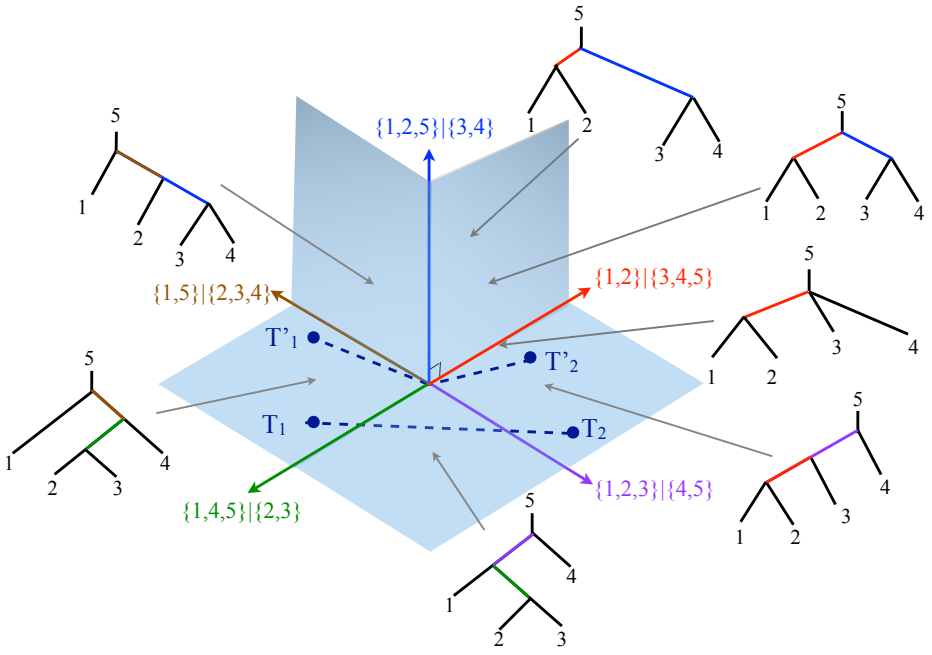

Figure 3: Five quadrants in $\mathcal{T}_5$. The upper right quadrant illustrates how trees with the same tree shape but different edge lengths lie at different coordinates in the quadrant. The dashed lines represent geodesics between two pairs of trees where T_i and T_i' have the same tree shape but different edges lengths. The geodesic T_1' to T_2' passes through the origin, while the geodesic T_1 to T_2 passes through an intermediate quadrant.

example, whether they are closed sets [6, Note 6.1.3.1]. We begin in subsection 4.1 by giving some examples to show that various properties of Euclidean convex hulls fail in CAT(0) complexes.

In subsection 4.2 we give our main result, an algorithm (using linear programming) to find convex hulls in any 2D CAT(0) complex with a single vertex O. Specifically, we prove:

Theorem 2. *There is a polynomial-time reduction from the problem of finding the convex hull of a finite set of points P in a 2D CAT(0) complex $\mathcal{K}$ with a single vertex O to linear programming. The resulting linear program has $O(n+m)$ variables and $O((n+m)^3)$ inequalities, where n is the number of cells in $\mathcal{K}$ and m is the number of points in P. For the special case of a cube complex this provides a polynomial-time convex hull algorithm.*

Our reduction uses the real-RAM model of computation. The idea of our algorithm is to first use the link graph to test if point O is in the convex hull and to identify the edges of the complex that intersect the convex hull at points other than O. Then we formulate the exact computation of the convex hull as a linear program whose variables represent the boundary points of the convex hull on the edges of the complex. There is a polynomial bound on the number of variables and inequalities of the linear program, but whether the linear program can be solved in polynomial time depends on bit complexity issues. There

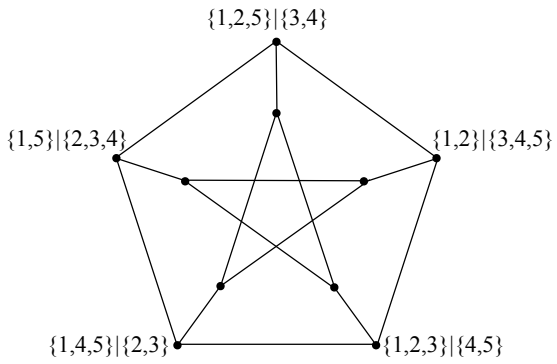

Figure 4: The Petersen graph, which is the link graph of the origin of $\mathcal{T}_5$. The graph edges correspond to tree shapes, and the graph vertices correspond to splits. The 5 quadrants in Figure 3 correspond to the outer 5-cycle in the Petersen graph.

are polynomial-time linear programming algorithms [33, 34], but their run-times depend on the number of bits in the input numbers. For cube complexes, which have angles of 90°, our linear program has coefficients with a polynomial number of bits and so our algorithm runs in polynomial time. However, more generally our algorithm must use the stronger real RAM model of computation in order to perform computations on the angles of the input CAT(0) complex, and we must resort to the simplex method for linear programming [19].

Our proof of Theorem 2 implies that the convex hull of a finite set of points in a single-vertex 2D CAT(0) complex is a closed set.

4.1 Counterexamples for Convex Hulls in CAT(0) complexes

In this section we give examples to show that the following properties of the convex hull of a set of points P in Euclidean space do not carry over to CAT(0) complexes.

1. Any point on the boundary of the convex hull of points in 2D is on a shortest path between two points of P.

2. In any dimensional space, the convex hull of three points is 2-dimensional.

3. Any point inside the convex hull can be written as a convex combination of points of P.

Our first example, shown in Figure 5, has three cells sharing an edge. Set P contains one point in each cell. The three shortest paths between pairs of points in P do not determine the convex hull. This shows that property 1 fails.

Furthermore, the example in Figure 6 shows that even in a single-vertex 2D CAT(0) rectangular complex, the convex hull of a set of points P can contain a point in a quadrant that is not entered by any shortest path between points of P. This example also shows

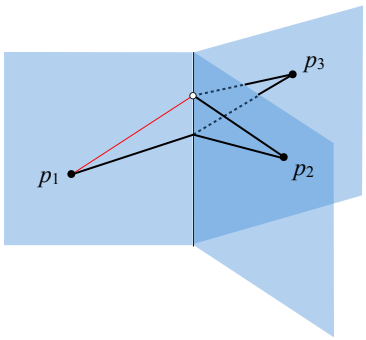

Figure 5: The shortest paths between p_1, p_2, p_3 (shown in black) do not determine the convex hull because the thin red line is on the boundary of the convex hull but not on any of the shortest paths.

that Carathéodory's property may fail, since point p is in the convex hull of the four points p_1, p_2, p_3, p_4 but not in the convex hull of any three of the points.

The example in Figure 7 shows that the convex hull of three points in a 3D CAT(0) complex may contain a 3D ball. This shows that property 2 fails.

Property 3 must be expressed more carefully for CAT(0) complexes because it is not clear what a convex combination of a set of points means except when the set has two points. If p and q are two points in a CAT(0) complex, then the points along the shortest path from p to q can be parameterized as $(1-t)p + tq$ for $t \in [0,1]$. Based on the definition of the convex hull, any point in the convex hull of a set of points P can be represented as a rooted binary tree with leaves labelled by points in P (with repetition allowed) and with the two child edges of each internal node v labelled by two numbers $(1-t_v)$ and t_v for $t_v \in [0,1]$, meaning that the point associated with v is this combination of the points represented by the child nodes. For example, see Figure 6.

One might hope that every point in the convex hull can be represented by a binary tree whose leaves are labelled by distinct elements of P. If this were true then we could verify that a point is in the convex hull by giving a binary tree with at most m leaves, and in the case of a cube complex, the problem of deciding membership in the convex hull would lie in NP. However, this hope is dashed by the example in Figure 7. Furthermore, the property may even fail for a 2D CAT(0) complex as we prove below for the example in Figure 8:

Lemma 1. *For the 2D CAT(0) complex shown in Figure 8, the point p cannot be represented by a binary tree with distinct leaves from* $\{p_1, p_2, p_3, p_4\}$.

Proof. In the link graph (shown in Figure 8(b)), the shortest path between p_3 and p_4 has length 180° so, by Property 1, the geodesic between p_3 and p_4 in the complex goes through the origin. Similarly, the geodesics between the following pairs also go through the origin: p_3 and p_2; p_3 and p_1; p_2 and p_4. The only pairs of points whose geodesic does not go through O are p_1, p_2, and p_1, p_4. The geodesic from p_1 to p_2 crosses axis 3 at point a. The geodesic from p_1 to p_4 crosses axis 5 at point c. Because p_1 is closer to axis 2, point c is closer to the origin than point a (see Figure 8(d)). The geodesic paths from p_3 to both a and c cross axis

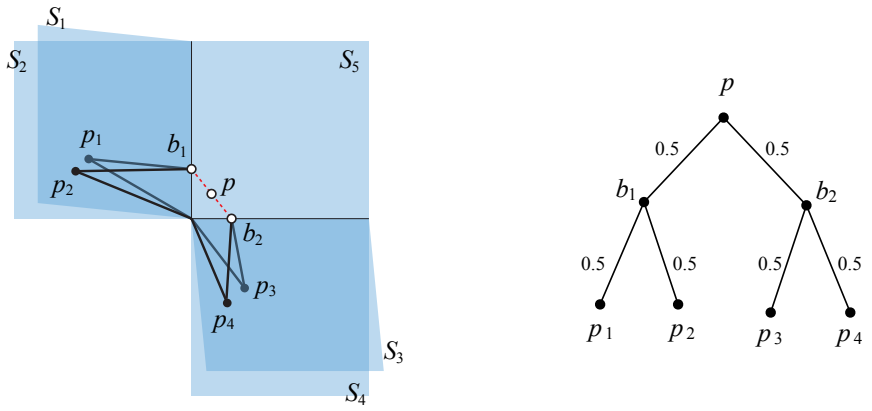

Figure 6: (left) A 2D CAT(0) space consisting of 5 quadrants. Quadrants S_1, S_2, S_5 share a vertical edge, and quadrants S_3, S_4, S_5 share a horizontal edge. Point $p_i, i = 1, 2, 3, 4$, lies in quadrant S_i. The convex hull of p_1, p_2, p_3, p_4 contains points—in particular p and the dashed red line—in a quadrant, S_5, that is not entered by any shortest path between points in P (shown as black lines). Note that 4 of the 6 shortest path between points of P go through the origin. (right) The tree representation of point p in terms of P.

4 and thus can be used to construct points in the cell bounded by axes 4 and 5 (where p lies). However, as shown in Figure 8(e), the geodesic path from p_3 to a crosses axis 4 further from the origin at point b, and therefore point p can only be constructed using points a and b, as shown by the binary tree in Figure 8(b). Therefore p can only be represented by a binary tree that repeats the leaf p_1. □

4.2 Convex Hull Algorithm for a Single-Vertex 2D CAT(0) Complex

In this section we prove Theorem 2 by reducing the convex hull problem for a single-vertex 2D CAT(0) complex $\mathcal{K}$ to linear programming via a polynomial-time reduction. Recall that P is the finite set of points whose convex hull we wish to find, and O is the single vertex of the complex.

4.2.1 Combinatorics of the Convex Hull

In this subsection we give algorithms to decide if O is in the convex hull and to identify the edges of $\mathcal{K}$ that contain points of the convex hull other than O. We will do this using the link graph $G = G_O$. Our algorithms take polynomial time in any model of computation where we can find shortest paths in the link graph and compare their lengths to π.

Testing if O is in the convex hull.

First consider the problem of testing whether O is in the convex hull. If $O \in P$ we are done, so assume that $O \notin P$. Then O is in $\mathrm{CH}(P)$ if and only if there is a path between two

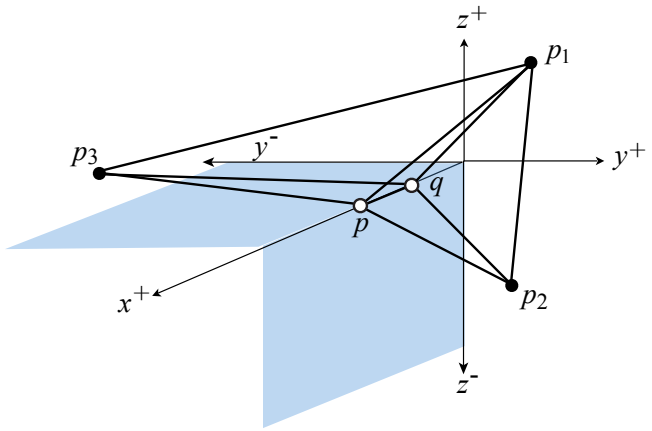

Figure 7: The convex hull of the three points p_1, p_2, p_3 in this 3D CAT(0) complex contains a 3D ball. Point p_1 lies in the $x^+y^+z^+$ octant, p_2 in the $x^+y^+z^-$ octant, and p_3 in the $x^+y^-z^+$ octant. The $x^+y^-z^-$ octant is missing. Point p is where the shortest path from p_2 to p_3 intersects the x^+ axis, and point q is where the triangle $p_1p_2p_3$ is pierced by the x^+ axis. The convex hull of P consists of the union of two simplices p_1p_3pq and p_1p_2pq.

points in $\mathrm{CH}(P) - O$ that goes through O, and Property 1 gives the conditions for a path to go through O. Thus:

Observation 1. *O is in $\mathrm{CH}(P)$ if and only if there are two points a and b in $\mathrm{CH}(P) - O$ such that the distance between a and b in the link graph is at least π.*

We test if O is in $\mathrm{CH}(P)$ as follows. If there are two points $p, q \in P$ whose distance in the link graph is at least π, then O is in $\mathrm{CH}(P)$. Otherwise, let $G[P]$ be the subset of the link graph that is the union of all shortest paths $\sigma_G(p, q),\ p, q \in P$. Note that $G[P]$ is not a subgraph of G in the usual sense because in general it includes portions of edges. Every point of $G[P]$ is the image of some point in $\mathrm{CH}(P)$. By the following lemma, it suffices to test if $G[P]$ has a cycle.

Lemma 2. *Suppose that no path between two points of P goes through O. Then $O \in \mathrm{CH}(P)$ if and only if $G[P]$ contains a cycle.*

Proof. Suppose $G[P]$ contains a cycle. Because the space is CAT(0), the cycle has length at least 2π, so it must contain two points a, b whose minimum distance in the cycle is π. Then the length of the shortest path between a and b in the link graph is π (otherwise we would have a cycle of length less than π). By Observation 1, O is in $\mathrm{CH}(P)$.

For the other direction, suppose $G[P]$ does not contain a cycle. $G[P]$ is connected, so it must be a tree. We claim that the leaves of the tree are points of P: If d is a point of $G[P]$ that is not in P then d is an internal point of some path $\sigma_G(p, q),\ p, q \in P$, so d has degree at least 2 in $G[P]$, so it is not a leaf.

Let a and b be points of $G[P]$. The path between a and b in $G[P]$ can be extended to a path between leaves of $G[P]$, and, since the leaves are in P, this path has length less than

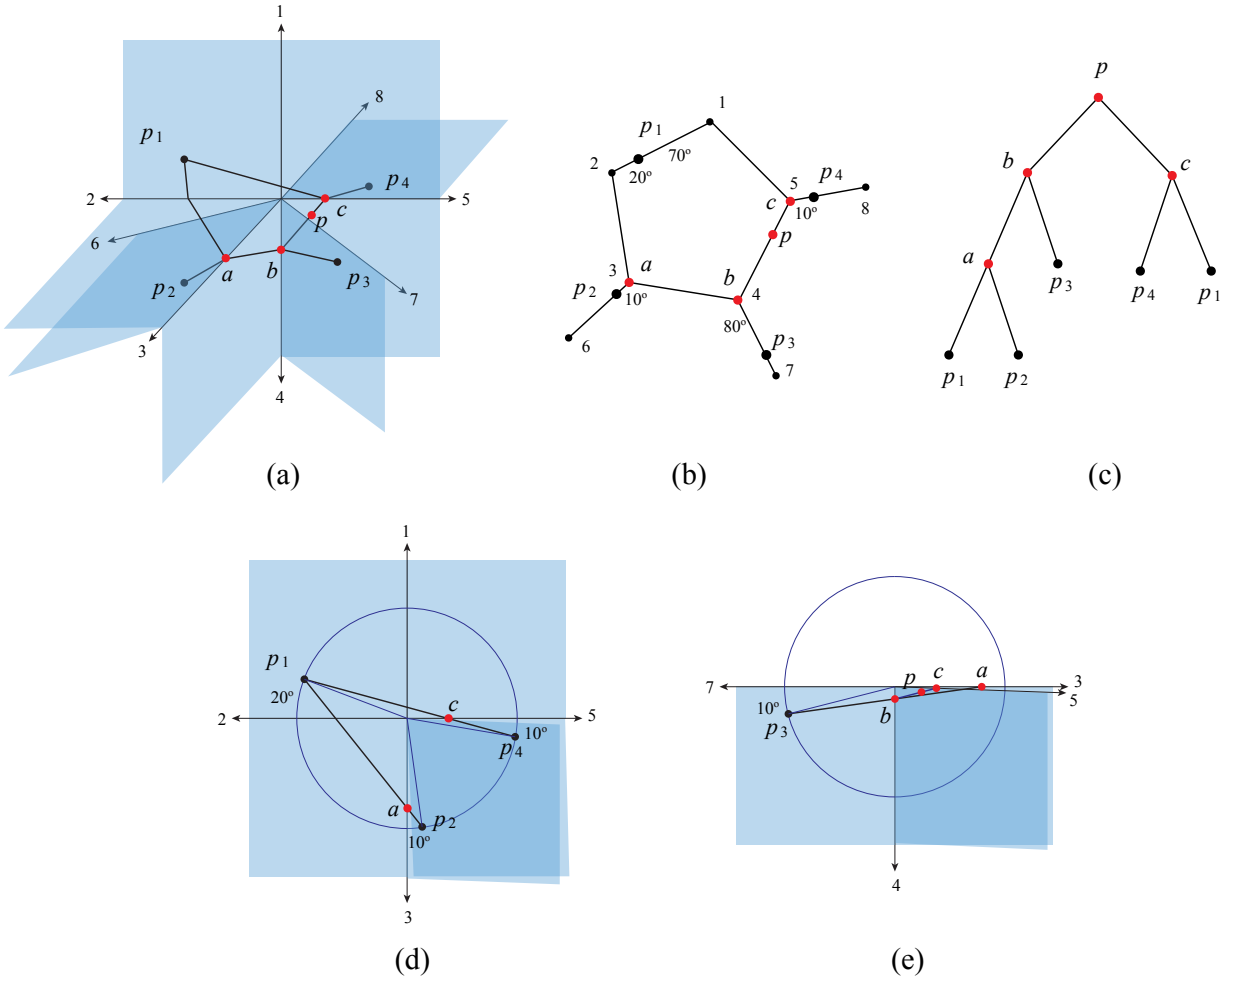

Figure 8: An example of a 2D CAT(0) complex where point p in $\mathrm{CH}(\{p_1, p_2, p_3, p_4\})$ cannot be represented as a binary tree with distinct leaves. Points p_1, p_2, p_3, p_4 are all distance 1 from the origin, with angles as specified in the link graph. (a) the complex (not to scale); (b) the link graph; (c) the binary tree representing p; (d) construction of points a and c; (e) construction of points b and p.

π. Thus there is a path between a and b in $G[P]$ of length less than π. This must be the shortest path between a and b in G. Thus, $G[P]$ is closed under taking shortest paths. We will never get two points at distance π or more. Therefore O is not in the convex hull. □

The algorithm to test if O is in $\mathrm{CH}(P)$ has a straight-forward implementation that runs in time $O(m(n+m))$ where n is the number of cells in $\mathcal{K}$ and m is the size of P. For each point $p \in P$, we do a depth-first search in the link graph to find paths to all the points of P within distance π of p. Note that a search to distance π will not find any cycles in the link graph, and therefore finds shortest paths from p. If some point $q \in P$ is not reached then we know that O is in $\mathrm{CH}(p)$. Otherwise, we construct $G[P]$ as the union of all these depth-first

search trees, and test if $G[P]$ contains a cycle. Each depth-first search takes time $O(n+m)$ so the total time is $O(m(n+m))$. Constructing and exploring $G[P]$ takes linear time.

Finding edges of $\mathcal{K}$ intersecting the convex hull.

We now show how to identify the edges of the complex $\mathcal{K}$ that contain points of $\mathrm{CH}(P) - O$. Equivalently, we will identify the vertices of the link graph that correspond to points in the convex hull. Our method applies whether or not O is inside the convex hull.

Let V be the set of vertices of the link graph. Then V corresponds to the set of edges of $\mathcal{K}$. Let B be the vertices of the link graph that correspond to points in the convex hull of P. We want to find B. The convex hull of P in the complex $\mathcal{K}$ is defined to be the closure of P under shortest paths. We observe that it suffices to consider shortest paths between points that are in P or lie on the edges of $\mathcal{K}$.

In the link graph this corresponds to taking shortest paths between points that lie in $P \cup V$. Furthermore, by Property 1 we should only take shortest paths in the link graph that have length less than π.

Therefore $P \cup B$ is the closure of the set P in the set $P \cup V$ under the operation of taking shortest paths of length less than π in the link graph.

Our algorithm finds this closure by building up a set $S \subseteq P \cup V$. Initially S will just be the input set of points P, and at the end of the algorithm, S will be the required set $P \cup B$. We will also keep a subset F of S that represents the "frontier" that we still need to explore from. Initially $F = S = P$.

The general step is to remove one element v from F. We then explore the part of the link graph within distance π from v. This can be done by a depth-first search in $O(n+m)$ time. As noted above, a search tree to distance π will find no cycles, and will therefore find shortest paths from v. We remove the part of the depth-first tree that is beyond the deepest point of S on each branch. Then for every vertex w of the link graph that is in the depth-first search tree, we check if w is already in S—if not then we add w to S and to F.

The size of S is bounded by $n+m$ where n is the number of cells in $\mathcal{K}$ and m is the size of P. Note that the amount of work we do for one element of F is $O(n+m)$. Thus the algorithm runs in time $O((n+m)^2)$.

The algorithm is correct because the final set S is closed under taking shortest paths of length less than π in the link graph. Each time we add a point to S we explicitly check all points of V within distance π from the point.

4.2.2 Finding the Convex Hull

In this section we give a polynomial-time algorithm to construct a linear program to find the convex hull of a finite point set P in a 2D CAT(0) complex $\mathcal{K}$ with a single vertex O. Our algorithm implies that the convex hull is closed in this case.

Suppose first that the origin O is inside the convex hull. The case where the origin is not inside the convex hull will be dealt with later. To find the convex hull it suffices to describe the intersection of the convex hull with each cell of the complex. Consider a cell C bounded by rays e and f incident to O. The part of the convex hull inside C is a convex polygon

determined by its vertices which consist of: point O; possibly some or all of the points of P inside C; and two points x_e and x_f on edges e and f, respectively, that are on the boundary of the convex hull. Note that we do not make any assumption about whether x_e and x_f are in the convex hull, because we are not making any assumption about whether the convex hull is open or closed. If we knew the points x_e and x_f, then we could easily compute the part of the convex hull inside C, since it is simply the Euclidean convex hull of x_e, x_f, O, and the points of P inside C.

There is an obvious iterative approach to finding the points x_ℓ where the boundary of the convex hull intersects each edge ℓ of the complex: Initialize H_0 to be the set P. For $i = 1, 2, \ldots$, initialize H_i to H_{i-1} and then take every pair of points from H_{i-1}, compute the shortest path σ between them, and add to H_i all the intersection points of σ with edges of the complex. (When two points lie on the same edge, we can discard the one closer to O.) We observe that this process is finite if and only if the convex hull is closed. As mentioned, our proof of Theorem 2 shows that the convex hull is closed and thus this iterative process is finite. In order for the iterative process to be an efficient algorithm we would need a bound on the number of iterations. We conjecture that there is a polynomial bound:

Conjecture 1. *$H_k = H_{k+1}$ for some polynomially bounded k.*

Lacking a proof of the conjecture, we will use linear programming to find the points where the boundary of the convex hull crosses the edges of the complex. Let B be the set of edges of the complex that have points other than O inside the convex hull. These are found as described in the previous section. Our algorithm and our notation will be simpler if the points of P all lie on edges of the complex. (In particular, the convex hull inside a cell C will be at most a triangle since no points of P will be internal to C.) We can achieve this by constructing an edge e_p from O through each point $p \in P$ (except the point O). Each such edge divides a cell in two. Point p is then represented in local coordinates by the distance along edge e_p from O to point p. We will use p to refer both to the point and to its local coordinate, i.e. its distance from O (in the same way that we refer to a point on the real line as a number).

Let $B' = B \cup \{e_p\}_{p \in P}$. For each $\ell \in B'$ our linear program will have a variable $x_\ell \in \mathbb{R}$ representing the distance from O to the point on ℓ that is on the boundary of the convex hull. Then $x_\ell > 0$. Note that, like p, x_ℓ refers both to a distance and a point.

Our inequalities are of two types. First, for any point $p \in P$ lying on an edge $\ell \in B'$ we include the inequality:

$$x_\ell \geq p \tag{1}$$

Inequalities of the second type will be determined by pairs of elements from the set B'. For any two edges e and f of B', such that the angle between e and f is $< \pi$, consider the shortest path σ between the corresponding points x_e and x_f. We will add a constraint for each edge ℓ of B' crossed by σ, expressing the fact that the convex hull includes the point where σ crosses ℓ. The constraint has the form $x_\ell \geq t$ where t is the distance from O to the

point where σ crosses ℓ. We will use t to refer to both the distance and to the point. We can express t in terms of known quantities. The set-up is illustrated in Figure 9. Note that there may be several polyhedral cells separating e and f, but we can lay them down to form a triangle.

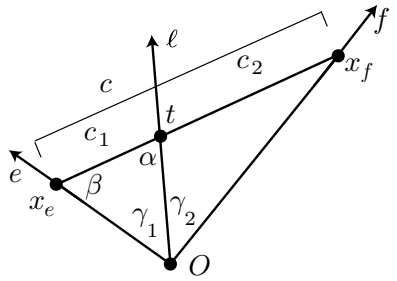

Figure 9: Expressing the intersection point t in terms of known quantities.

Let γ_1 be the angle between e and ℓ, and let γ_2 be the angle between ℓ and f. Then we have:

Claim 1.

$$t = \frac{x_e x_f \sin(\gamma_1 + \gamma_2)}{x_e \sin\gamma_1 + x_f \sin\gamma_2}$$

Proof. We apply the sine law several times. Let α and β be the angles opposite sides Ox_e and Ot, respectively, in the triangle Ox_et, and let c be the distance from x_e to x_f. By the sine law

$$t = \frac{x_e \sin\beta}{\sin\alpha}, \quad \sin\beta = \frac{x_f \sin(\gamma_1 + \gamma_2)}{c}.$$

Combining these we get

$$t = \frac{x_e x_f \sin(\gamma_1 + \gamma_2)}{c \sin\alpha}.$$

Let c_1 and c_2 be the parts of c opposite γ_1 and γ_2, respectively. By the sine law, we also have $c_1 \sin\alpha = x_e \sin\gamma_1$ and $c_2 \sin(\pi - \alpha) = c_2 \sin\alpha = x_f \sin\gamma_2$. So $c \sin\alpha = (c_1 + c_2)\sin\alpha = x_e \sin\gamma_1 + x_f \sin\gamma_2$. Substituting into the above expression for t gives the desired formula. □

Using the above claim, the constraint $x_\ell \geq t$ becomes

$$x_\ell \geq \frac{x_e x_f \sin(\gamma_1 + \gamma_2)}{x_e \sin\gamma_1 + x_f \sin\gamma_2}$$

This is not a linear inequality, but substituting $y_\ell = \frac{1}{x_\ell}$ yields

$$y_\ell \leq y_f \frac{\sin\gamma_1}{\sin(\gamma_1+\gamma_2)} + y_e \frac{\sin\gamma_2}{\sin(\gamma_1+\gamma_2)} \tag{2}$$

Since the γ_i's are constant, this is a linear inequality.
The inequalities (1) for point $p \in P$ on edge ℓ become

$$y_\ell \leq \frac{1}{p} \tag{3}$$

We claim that maximizing $\sum y_\ell$ subject to the inequalities (2), (3) and $y_\ell \geq 0$ gives the convex hull of P. The argument is as follows. First note that the closure of $\mathrm{CH}(P)$ provides a feasible solution to the linear system, because the convex hull satisfies the constraints $x_\ell \geq t$ which we used to construct our inequalities. Denote this solution by $y_\ell^{\mathrm{CH}}, \ell \in L$. Next, note that any other solution $y'_\ell, \ell \in L$, has $y'_\ell \leq y_\ell^{\mathrm{CH}}$ for all $\ell \in L$—in other words, any other solution includes $\mathrm{CH}(P)$. This is because the points of P are included, and the inequalities enforce closure under shortest paths. Therefore, the solution that maximizes $\sum y_\ell$ is the convex hull of P.

Thus we have reduced the problem of finding the convex hull (when O is in the convex hull) to linear programming.

We now deal with the case where the origin O is not inside the convex hull. In this case, by Lemma 2, the subgraph of the link graph corresponding to the convex hull is a tree, and it seems even more plausible that an efficient iterative approach can be used to find the convex hull. However, we leave this as an open question, and give a linear programming approach like the one above.

As before, we will construct a new edge through every point of P, and let B' be the set of edges (including the newly constructed ones) that have points other than O inside the convex hull. For each $\ell \in B'$ we will make two variables, $x_\ell^{\min}$ and $x_\ell^{\max}$ in $\mathbb{R}$ representing the minimum and maximum points on ℓ that are on the boundary of the convex hull. To find the convex hull, it suffices to find the values of these variables.

We want to ensure that $x_\ell^{\max}$ is larger than any point of P and any point at which a geodesic between $x_e^{\max}$ and $x_f^{\max}$, for any $e, f \in B'$, crosses edge ℓ. Similarly, we want to ensure that $x_\ell^{\min}$ is smaller than any point of P and any point at which a geodesic between $x_e^{\min}$ and $x_f^{\min}$, for any $e, f \in B'$, crosses edge ℓ. Using the same notation and set-up as above with edges e and f, and using the inverse variables $y_\ell^{\min} = \frac{1}{x_\ell^{\min}}$ and $y_\ell^{\max} = \frac{1}{x_\ell^{\max}}$ the inequalities corresponding to (2) are:

$$y_\ell^{\max} \leq y_f^{\max} \frac{\sin\gamma_1}{\sin(\gamma_1+\gamma_2)} + y_e^{\max} \frac{\sin\gamma_2}{\sin(\gamma_1+\gamma_2)}$$
$$y_\ell^{\min} \geq y_f^{\min} \frac{\sin\gamma_1}{\sin(\gamma_1+\gamma_2)} + y_e^{\min} \frac{\sin\gamma_2}{\sin(\gamma_1+\gamma_2)}$$

The inequalities corresponding to (3) are:

$$y_\ell^{\max} \leq \frac{1}{p} \leq y_\ell^{\min}$$

If we maximize the objective function $\sum(y_\ell^{\max} - y_\ell^{\min})$ subject to the above inequalities and $y_\ell^{\min} \geq y_\ell^{\max} \geq 0$ then, by a similar argument to the one above, this gives the convex hull of P. Thus we have reduced the problem of finding the convex hull to linear programming.

We now discuss the running time of the algorithm. We will concentrate on the case where O is in the convex hull—the other case is similar. Recall that n is the number of cells in the complex and m is the number of points in P. At the beginning of the algorithm we test if O is in the convex hull, and find the set B of edges of the complex that contain points of the convex hull other than O. This takes $O((n+m)^2)$ time as discussed in the previous section. The set B' has size $O(m+n)$ because it includes an edge of the complex through every point of P. The linear program has $O(n+m)$ variables. The number of inequalities is $O((n+m)^3)$ since we consider each pair of elements, e, f from B', and add an inequality for each edge of the complex crossed by the shortest path from e to f. We can construct the linear program in polynomial time assuming a real RAM model of computation that allows us to do computations on the input angles (including computing the sine of angles).

The special case of a cube complex

In the special case of a cube complex, it is more natural to give each input point using x- and y-coordinates relative to the quadrant containing the point. In this case, we claim that all the low-level computations described above can be performed in polynomial time when measuring bit complexity. We will not construct new edges through points of P since that introduces new angles. Our variables are x_e for $e \in B$, and we add constraints for shortest paths between pairs of points in $P \cup \{x_e\}_{e \in B}$.

We give a few more details for the computation of t in Figure 9 in this case. Figure 9 shows a path between x_e and x_f crossing an edge ℓ at t. In the current setting, ℓ will be an edge of B, and x_e and x_f may be variables or input points. If x_e and x_f correspond to input points, then t is just the point where a line between two known points crosses an axis. Then the right-hand-side of the corresponding constraint (2) is a constant whose bit complexity is polynomially bounded in terms of the input bit complexity. The case when both x_e and x_f are variables cannot arise because they would be distance π apart in the link graph. Thus the only case we must take care of is when x_e corresponds to an input point and x_f is a variable (or vice versa). The situation is shown in Figure 10, with x_e being an input point p with coordinates (h, v) as shown. Then $t = vx_f/(x_f + h)$, so constraint (2) becomes $y_\ell \leq (hy_f + 1)/v$. Thus the coefficients in our linear constraints are rationals whose bit complexity is polynomially bounded in terms of the input bit complexity. This means that we can use polynomial-time linear programming algorithms [33, 34] to find the value of x_e for each edge $e \in B$.

Finally, as noted at the beginning of this section, we can use the values x_e to compute the convex hull in each quadrant C. We repeat the argument here. Suppose C is bounded by edges e and f. Then the convex hull inside C is the Euclidean convex hull of points x_e,

x_f, O, and the points of P that lie inside C, and since Euclidean (planar) convex hulls can be computed in polynomial time [12], this allows us to compute the convex hull in C, and thus the whole convex hull, in polynomial time.

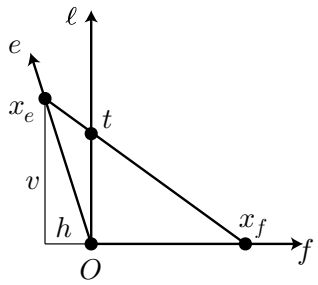

Figure 10: Expressing the intersection point t in terms of known quantities in the case of a cube complex.

This completes the proof of Theorem 2.

Finally, we observe that since the extreme points of the convex hull can be expressed via linear inequalities, therefore the convex hull of a finite point-set in a single-vertex 2D CAT(0) complex is closed.

5 Shortest Paths

In this section we explore the possibilities and limitations of using the shortest path map to solve the single-source shortest path problem in a 2D CAT(0) complex. The input is a 2D CAT(0) complex, $\mathcal{K}$, composed of n triangles, and a "source" point s in $\mathcal{K}$. We denote the shortest path from s to t by $\sigma(s,t)$.

We first show that the shortest path map may have exponential size in general. This contrasts with the fact that the shortest path map has size $O(n^2)$ in the two special cases where the single-source shortest path problem is known to be efficiently solvable: when the complex is a topological 2-manifold with boundary, which we will call a 2-manifold for short [39]; and when the complex is rectangular [18].

We then show that for any 2D CAT(0) complex there is a structure called the "last step shortest path map" that coarsens the shortest path map, has size $O(n)$, and allows us to find the shortest path $\sigma(s,t)$ to a given target point t in time proportional to the number of triangles and edges traversed by the path. Although we do not know how to find the last step shortest path map in polynomial time for general 2D CAT(0) complexes, we can obtain it from the shortest path map.

From this, we obtain efficient algorithms for the single-source shortest path problem in 2D CAT(0) complexes that are 2-manifold or rectangular. Both cases had been previously solved, but the techniques used in the two cases were quite different. Our approach is the same in both cases and opens up the possibility of solving other cases. We need $O(n^2)$ preprocessing time and space to construct a structure that uses $O(n)$ space and allows us to

find the shortest path $\sigma(s,t)$ to a given target point t in time proportional to the number of triangles and edges traversed by the path. This improves the bounds for 2-manifolds, and is competitive for rectangular complexes (where the published bounds are for two-point queries, so the situation is not quite comparable).

5.1 The Shortest Path Map

Typically in a shortest path problem, the difficulty is to decide which of multiple geodesic (or locally shortest) paths to the destination is shortest. This is the case, for example, for shortest paths in a planar polygon with holes, or for shortest paths on a terrain, and is a reason to use a Dijkstra-like approach that explores paths to all target points in order of distance. For shortest paths on a terrain, Chen and Han [15] provided an alternative that uses a Breadth-First-Search (BFS) combined with a clever pruning when two paths reach the same target point.

When geodesic paths are unique, however, it is enough to explore all geodesic paths, and there is no need to explore paths in order of distance or in BFS order. This is the case, for example, for shortest paths in a polygon, where the "funnel" algorithm [26,29] achieves $O(n)$ processing time and storage, and $O(\log n)$ query time (plus output size to produce the actual path). Similarly, in CAT(0) spaces, the uniqueness of geodesic paths means we can obtain a correct algorithm by simply exploring all geodesic paths without any ordering constraints.

A basic approach to the single source shortest path problem is to compute the whole *shortest path map* from s. The *shortest path map* partitions the space into regions in which all points have shortest paths from s that have the same *combinatorial type*. Specialized to 2D CAT(0) complexes, two shortest paths have the same *combinatorial type* if they cross the same sequence of edges, vertices, and faces.

For a vertex v in a 2D CAT(0) complex, we define the *ruffle* of v to be the set of points p in the complex such that the shortest path from s to p goes through v. See Figure 11 for an example in the case of a rectangular complex. The points of the ruffle of v in a small neighbourhood of v can be identified from the link graph of v together with the incoming ray which is the last segment of the shortest path $\sigma(s,v)$. In particular, the points of v's ruffle close to v are those points p for which the segment vp makes an angle of at least π with the incoming ray. Using the link graph, the boundary rays of the ruffle of v can be identified in time proportional to the number of faces incident to v.

Consider one region of the shortest path map, and the set, C, of shortest paths to points in the region. The paths in C all go through the same sequence, S_C, of faces and edges and vertices. Let v be the last vertex in the sequence S_C (possibly $v = s$). There is a unique geodesic path from s to v, and all the paths of C traverse this same path from s to v. After that, the points of the paths of C all lie in the ruffle of v. Since the paths traverse the same sequence of edges and faces they can be laid out in the plane to form a cone with apex v. See Figure 12. Observe that the boundary rays of the cone may or may not lie in the set C. If the boundary of the cone is the boundary of the ruffle of v then it is included in C; but if the boundary of the cone is determined by another vertex, then beyond that vertex, the boundary is not included. Note however, that the boundary ray is a shortest path—just not

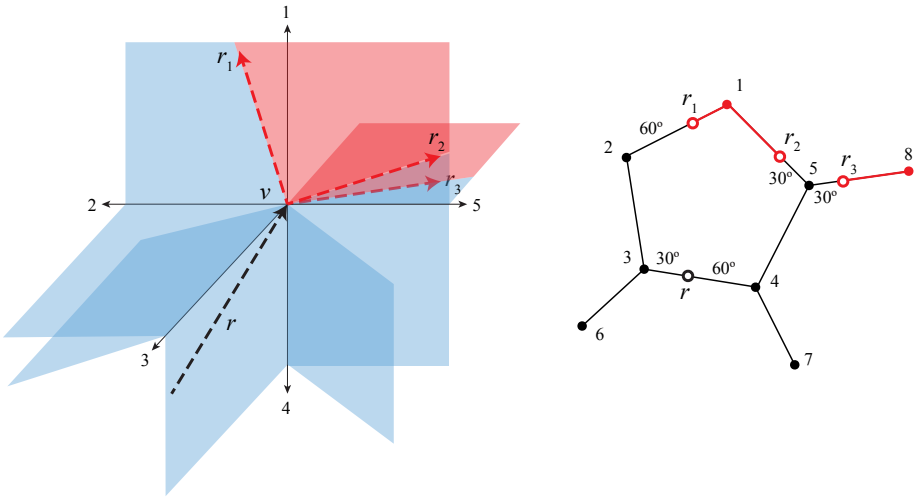

Figure 11: The ruffle (in red) of vertex $v \in \mathcal{K}$ with respect to incoming ray r, shown in $\mathcal{K}$ (left) and in the link graph G_v (right). The boundary rays of the ruffle are r_1, r_2, and r_3.

of the same combinatorial type since it goes through another vertex.

5.1.1 Computing the shortest path map

We will show that if the shortest path map has M regions, then it can be computed in time $O(M)$. Regions of the shortest path map may have dimension 0, 1, or 2. Each 2-dimensional region of the shortest path map is bounded by: two boundary rays; a vertex or a segment of an edge through which shortest paths enter the region; and one or two segments of edges and possibly a vertex through which shortest paths exit the region. See Figure 12. With each region, we will store its boundary rays and vertices/segments. Each vertex of the complex is a 0-dimensional region of the shortest path map. An edge may form a 1-dimensional region of the shortest path map (for example any edge (v, w) inside the ruffle of v).

The algorithm builds the regions of the shortest path map working outwards from s. In general, we will have a set of vertices and segments (portions of edges) that form the "frontier" of the known regions, and at each step of the algorithm, we will advance the known regions beyond one frontier vertex/segment.

The algorithm is initialized as follows. Assume that s is a vertex of the complex (if necessary, by triangulating the face containing s). Each edge incident to f becomes a region of the shortest path map. Each face f incident to s becomes a region of the shortest path map with the two edges of f that are incident to s as its boundary rays. The two vertices of f different from s enter the frontier, along with the edge of f not incident to s.

At each step of the algorithm we take one vertex or segment out of the frontier set and we find all the regions for which shortest paths enter through this vertex or segment.

Consider first the case of removing segment s from the frontier. We wish to find the regions of the shortest path map for which shortest paths enter through segment s. If

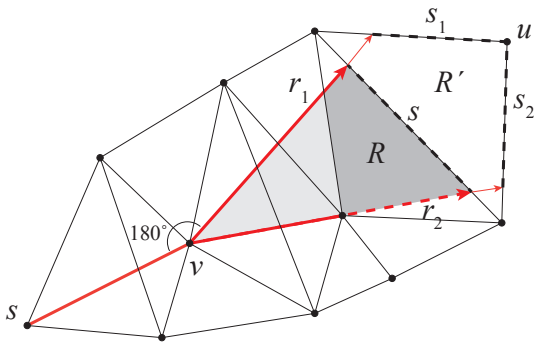

Figure 12: The structure of shortest paths to one region R (shown darkly shaded) of the shortest path map. The set C of shortest paths to points in the region forms a path $\sigma(s, v)$ together with a cone (lightly shaded) with apex v bounded by rays r_1 and r_2. Region R is closed on the r_1 boundary and open on the r_2 boundary. Shortest paths exit R through segment s. The figure shows one region of the shortest path map beyond s that has two segments s_1 and s_2 and a vertex u through which shortest paths exit. Note that the angles of the triangles incident to v are not drawn accurately since they sum to more than 2π.

segment s lies in edge e, then the faces containing the new regions are those incident to e, not including the face from which shortest paths arrive at s. (See segment s and region R' in Figure 12 for example.) Each such region R' gives rise to one or two segments and possibly a vertex through which shortest paths exit the region. We add these segments and vertex to the frontier. In case there is a vertex, u, (such as in Figure 12) we must find the shortest path to the vertex. This can be done by placing the boundary rays of R' in the plane, computing their point of intersection, p, and constructing the ray from p to u. Note that we do not need to know the sequence of faces traversed by shortest paths to region R'—local information suffices. This provides us with the shortest path to u and also the boundary rays of the segments incident to u.

We next consider the case where a vertex v is removed from the frontier. We must find the regions of the shortest path map for which shortest paths enter through vertex v. These lie in the ruffle of v. Knowing the shortest path $\sigma(s, v)$, we can search the link graph G_v of v to find all the boundary rays of the ruffle of v. Any edge incident to v that lies in the ruffle forms a 1-dimensional region of the shortest path map, and we add its other endpoint to the frontier. For each face f incident to v, we can identify the region of the shortest path map that lies in face f and interior to the ruffle of v. We can also identify the segments and vertices through which shortest paths exit the new region, and add these to the frontier.

This completes the high-level description of the algorithm. We spend constant time per region of the shortest path map, plus $O(n)$ time to search the faces incident to each vertex, for a total of $O(M)$.

If we want to use the shortest path map to answer shortest path queries, we also need a way to locate, given a target point t that lies in face f, which region of the shortest path map contains t. This necessitates building a search structure for the shortest path regions

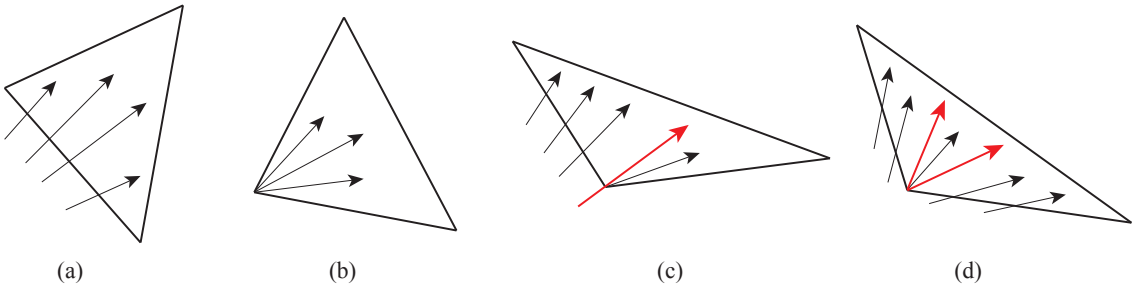

Figure 13: Shortest paths may enter a face through: (a) one edge (type E); (b) one vertex (type V); (c) one edge and an incident vertex (type EV); or (d) two edges and their common vertex (type EVE).

that face f is partitioned into, which takes more time and space. (Results of Mount [43] might might give a solution better than the obvious one for this.) We will not pursue this solution because we will present an alternative solution in Section 5.2.

5.1.2 Properties of the shortest path map

For our remaining results, we need some properties of shortest paths in a 2D CAT(0) complex.

Lemma 3. *Let e be an edge of a 2D CAT(0) complex. Either all the shortest paths to internal points of e travel along e, or they all reach e from one incident face.*

Proof. If the shortest path to some internal point of edge e travels along e, then so do the shortest paths to all internal points of e.

Otherwise consider the (finitely many) combinatorial types of shortest paths to points of e, and let $C_1, C_2, \ldots, C_k$ be the corresponding sets of shortest paths, ordered according to the order of points along e. We will prove that paths in all the C_i's arrive at points of e from the same incident face. For otherwise, there would be some C_i and C_{i+1} that arrive from different incident faces. The boundary ray between C_i and C_{i+1}, must be part of one or the other, say C_{i+1}. But observe that when C_i is laid out in the plane, the boundary ray of its cone (on the side of C_{i+1}) is still a shortest path, and still arrives at e from the same incident face as C_i does. But this contradicts C_{i+1} arriving from a different face. □

We next characterize how shortest paths can enter a face (a triangle) of the complex. See Figure 13.

Lemma 4. *Shortest paths enter a triangular face either through one edge, or one vertex, or one edge and an incident vertex, or two edges and their common vertex.*

Proof. We cannot have shortest paths entering a face from all three edges, nor from an edge and the opposite vertex, otherwise we would have shortest paths to two points on the same edge arriving from different faces, in contradiction to Lemma 3. □

5.1.3 Size of the shortest path map

A boundary ray between adjacent regions of the shortest path map starts out as a boundary ray of the ruffle of some vertex. By Lemma 4, each face originates at most two such rays. In a general 2D CAT(0) complex, such a ray can bifurcate into two or more branches when it hits an edge that is incident to more than two faces. There is one branch for each new incident face. See Figure 15(a) for an example. The collection of all branches that originate from one boundary ray of a ruffle is called a *boundary tree*. Observe that it is a tree—no two branches can intersect because geodesic paths are unique. There are $O(n)$ boundary trees because each face originates at most two boundary trees. If the complex is a 2-manifold (i.e., every edge is in at most two faces) then no bifurcations can occur, so each boundary tree consists of only one branch, which implies that the size of the shortest path map is $O(n^2)$. This was proved by Maftuleac [39] (where 2-manifold complexes are called "planar"), but we include a proof because we wish to observe a generalization.

Lemma 5 ([39])**.** *In a 2D CAT(0) complex that is a 2-manifold the size of the shortest path map is $O(n^2)$.*

Proof. As noted above, every boundary tree consists of only one branch, or ray. If such a ray entered a face twice then the second entry would not be a shortest path, since we could short-cut across the face from the first entry. Therefore no ray enters a face twice, and the number of boundary tree branches cutting any face is $O(n)$. Then the number of regions of the shortest path map within one face is $O(n)$ and the overall number of regions is $O(n^2)$. □

In a general 2D CAT(0) complex it may happen that no two branches of one boundary tree cross the same face, which implies that the shortest path map still has size $O(n^2)$. We prove that this is the case for 2D CAT(0) rectangular complexes:

Lemma 6. *In a 2D CAT(0) rectangular complex, no two branches of one boundary tree can enter the same face, and from this it follows that the shortest path map has size $O(n^2)$.*

Proof. Suppose, by contradiction, that two branches r_1 and r_2 of the same boundary tree enter a common face. Let f be the first face they both enter. Edges of a rectangular complex lie in two classes, which we call "horizontal" and "vertical", such that edges in one class are incident only to edges from the other class. Observe that if a ray makes an angle of α with some horizontal edge, then it makes the same angle α with every horizontal edge that it crosses in the rectangular complex, and it makes the same angle $\pi/2 - \alpha$ with every vertical edge that it crosses. See Figure 14. Consequently, r_1 and r_2 are parallel in f. Consider a line segment b joining a point of r_1 and a point of r_2 in f, and consider the shortest paths that arrive at points of b. The last segments of all these shortest paths must be parallel. This contradicts the fact that the set of shortest paths corresponding to any region of the shortest path map form a cone of rays.

Therefore the branches of one boundary tree enter a face at most once. Since there are $O(n)$ boundary trees, this means that the number of boundary tree branches cutting any face is $O(n)$. Then the number of regions of the shortest path map within one face is $O(n)$ and the overall number of regions is $O(n^2)$. □

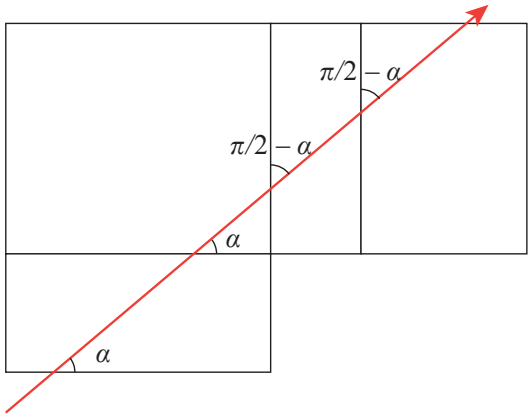

Figure 14: If a ray makes an angle of α with some horizontal edge in a 2D CAT(0) rectangular complex, then it makes the same angle α with every horizontal edge that it crosses, and it makes an angle $\pi/2 - \alpha$ with every vertical edge that it crosses.

In a general 2D CAT(0) complex, two branches of one boundary tree may cross the same face—see Figure 15(a) for an example—and the size of the shortest path map may grow exponentially:

Lemma 7. *The size of the shortest path map of a 2D CAT(0) complex may be exponential in n, the number of faces.*

Proof. Figure 15(b), (c) show how one boundary ray of a ruffle can bifurcate into two branches which then enter the same face g_1. Figure 15(d) shows how this process can be repeated. With each addition of three faces, f_i', f_i'', and g_i, the number of branches doubles. Thus after adding $3n$ faces, the number of branches is 2^n. □

Note that an exponential size shortest path map does not preclude polynomial time algorithms for computing shortest paths. In the tree space and its generalization, orthant space, the shortest path map, and indeed the number of regions in a face, can have exponential size [40, 45], but there is still a polynomial time algorithm for computing geodesics in these spaces [40, 46].

5.2 The Last Step Shortest Path Map

Although the shortest path map for single-source shortest paths in a 2D CAT(0) complex may have exponential size, there is a structure, called the "last step shortest path map," that has linear size and can be used to find the shortest path to a queried target point in time proportional to the combinatorial size of the path (i.e., the number of faces, edges, and vertices traversed by the path).

The *last step shortest path map*, first introduced in [20], partitions the space into regions where points p and q are in the same region if the shortest paths $\sigma(s, p)$ and $\sigma(s, q)$ have

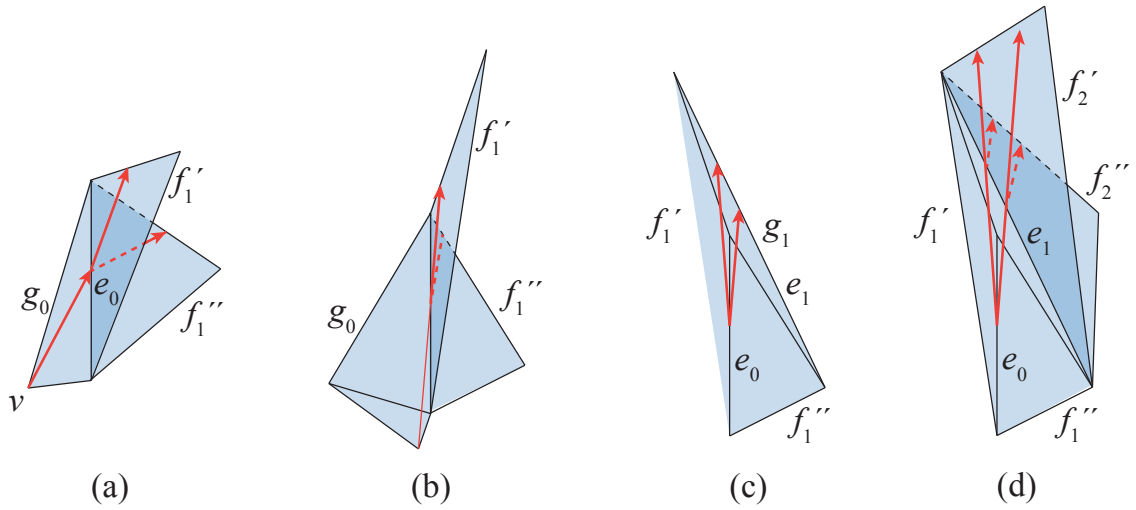

Figure 15: (a) A boundary ray of a ruffle (shown in red with arrows) originates from vertex v in face g_0, and bifurcates when it reaches edge e_0, branching into two rays, one in face f_1' and one in face f_1''. (b) The same situation but with sharper angles. (c) The two resulting branches enter face g_1 that is incident to f_1' and f_1'', and arrive at edge e_1. Note that (b) and (c) show opposite sides of face f_1'. (d) Two more faces f_2' and f_2'' are incident to edge e_1, so the two branches bifurcate into a total of four branches. In the next iteration, the four branches will enter a face g_2 incident to f_2' and f_2''. The process can be continued, and the number of rays doubles each time we add three faces.

the same last vertex, edge, or face, i.e., the combinatorial type of the two paths matches on the last element. Thus, the last step shortest path map is a coarsening of the shortest path map.

In the last step shortest path map, each edge has a unique incoming edge or face (by Lemma 3). By Lemma 4 each face is divided into at most three regions, as shown in Figure 13. Thus for any 2D CAT(0) complex the last step shortest path map has size $O(n)$. For the purpose of answering shortest path queries, we store with each region the last vertex, edge, or face with which shortest paths enter the region. We call this the *incoming* information for the region.

5.2.1 Answering shortest path queries using the last step shortest path map

We show that the last step shortest path map, together with the incoming information described above, is sufficient to recover the path from s to any point t in time proportional to the number of faces on the path. A query point t is given as a vertex, or a point on an edge, or a point (in local coordinates) in a face.

We first test whether a shortest path reaches t along an edge—this happens if and only if t is a vertex or a point on an edge and the incoming information attached to the vertex or edge is an edge, say the edge from u to v. In this case, we replace t by u and recurse. Otherwise,

t lies in a face, edge or vertex and the incoming information tells us that a shortest path reaches t through a face, say f. Refer to Figure 13. If f is of type V, we replace t by the incoming vertex of f and recurse. If f is of type VE or type EVE we locate t relative to the rays in f. From this we can tell if the shortest path to t goes through a vertex of f or not. If it does, then we replace t by that vertex and recurse.

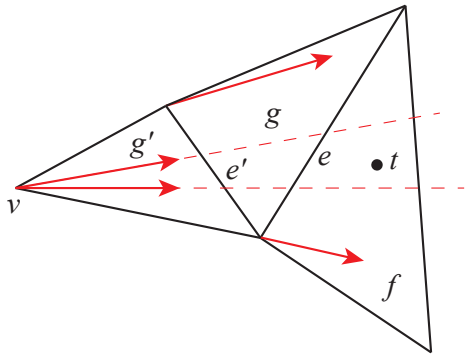

Figure 16: Finding the shortest path from s to point t in face f. In this example, f is of type VE. Testing the ray of f, we find that the shortest path to t enters from edge e which has incoming face g of type VE. Testing the rays of g, we find that the shortest path to t enters from edge e' which has incoming face g' of type EVE. Finally, testing the rays of g' we find that the shortest path to t comes from vertex v. We recursively find the shortest path to v.

We are left with the case where the shortest path to t enters face f through some edge, say edge e. Let g be the incoming face for edge e. We place f in the plane and attach triangle g to edge e. The placement of f is arbitrary, but then t and g are fixed. Now we enter the main loop of the algorithm (see Figure 16): If g is of type V, we replace t by the incoming vertex of g and recurse. If g is of type VE or type EVE we locate t relative to the rays in g (although t is not in g we just extend the rays to do the test). From this we can tell if the shortest path to t goes through a vertex of g or not. If it does, then we replace t by that vertex and recurse. Otherwise the shortest path to t enters g through an edge, and we repeat with the incoming face of that edge.

This algorithm finds the shortest path from s to t in time proportional to the number of triangles and edges on the path. In the worst case this is $O(n)$.

5.2.2 Computing the last step shortest path map

We do not know how the compute the last step shortest path map in polynomial time. More broadly, we do not know of a polynomial-time algorithm to compute shortest paths in a 2D CAT(0) complex. On the other hand, the problem does not seem to be amenable to NP-hardness proofs like the ones for shortest paths in 3D Euclidean space with polyhedral obstacles [11], or for shortest paths that visit a sequence of non-convex polygons in the plane [20]. Furthermore, we have the example of orthant spaces as CAT(0) complexes with

exponential shortest path maps, but a polynomial time algorithm for computing shortest paths [40].

It is tempting to think that the last step shortest path map can be computed in a straightforward way by propagating incoming information outward from the source. The trouble with this approach is that faces of type *EVE* need incoming information from two edges. This can result in dependencies that form a cycle, with each edge/face waiting for incoming information from some other face/edge. See Figure 17 for an example.

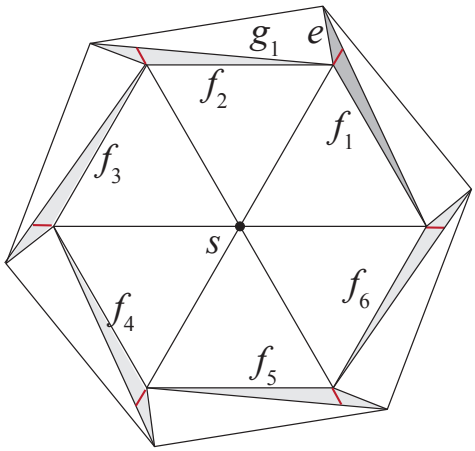

Figure 17: A cycle of incoming information in a 2D CAT(0) complex that lies in the plane. Face f_1 (darkly shaded) is of type *EVE* with incoming edge e, which has incoming face g_1, which depends on incoming information from face f_2. Similarly, each face f_i (lightly shaded) depends on incoming information from face f_{i+1}, and f_6 depends on incoming information from face f_1, which creates a cycle.

We can compute the last step shortest path map from the shortest path map in time $O(M)$. For each edge, we can identify the incoming edge or face from any of the shortest path regions containing portions of the edge (by Lemma 3 these all give the same information). Since we have the shortest path to each vertex v, we can recover or recompute the boundary rays of the ruffle of v, which gives us the type (*E*, *V*, *EV*, or *EVE*) of each face incident to v, and the incoming information for the face.

We summarize the implications for special cases of the single-source shortest path problem in 2D CAT(0) complexes:

Lemma 8. *For a 2D CAT(0) complex that is a 2-manifold or is rectangular, we can solve the single-source shortest path problem using $O(n^2)$ time and space to produce a structure (the last step shortest path map) of size $O(n)$ that allows us to answer shortest path queries in time proportional to the number of triangles and edges on the path.*

6 Conclusions

We have given an algorithm for computing the convex hull of a set of points in a 2D CAT(0) polyhedral complex with a single vertex. Our algorithm relies on linear programming. We leave it as an open question to find a combinatorial algorithm, or prove that a simple iterative approach takes polynomial time (Conjecture 1).

We do not see how to extend our linear programming solution to 2D CAT(0) complexes with more than one vertex, nor to single-vertex higher dimensional CAT(0) complexes. In 2D, the boundary between maximal cells is an edge, and the convex hull intersects the edge in an interval. In higher dimensions, the boundary between two maximal cells has dimension 2 or greater, and it is not clear that the intersection of the convex hull with a boundary face is even a polytope. Rather than computing the convex hull explicitly, it might be easier to find an algorithm that tests whether a point is in the convex hull or not. This would be sufficient for most applications, including computing a geometric centre by peeling convex hulls.

For the single-source shortest path problem in a 2D CAT(0) complex, we have shown that the shortest path map may have exponential size, and that the last step shortest path map is a better alternative. It is an open question to compute the last step shortest path map in polynomial time. Alternatively, one might try to prove that the shortest path problem is NP-hard for 2D CAT(0) complexes.

7 Acknowledgements

The authors thank Sean Skwerer for the example showing convex hulls of 3 points can be 3-dimensional (Figure 7), and Aasa Feragen, Steve Marron, Ezra Miller, Vinayak Pathak, Scott Provan, and Sean Skwerer for helpful discussions about convex hulls in tree space. We are extremely grateful to anonymous reviewers for helping us catch an error in an earlier version.

MO acknowledges the support of the Fields Institute. Research supported by NSERC, the Natural Sciences and Engineering Research Council of Canada.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ, 1993.

[2] F. Ardila, M. Owen, and S. Sullivant. Geodesics in CAT(0) cubical complexes. *Advances in Applied Mathematics* 48(1):142–163, 2012.

[3] M. Bacak. Computing medians and means in Hadamard spaces. *SIAM Journal on Optimization* 24(3):1542–1566, 2014.

[4] M. Bacak. *Convex analysis and optimization in Hadamard spaces*, vol. 22. Walter de Gruyter GmbH & Co KG, 2014.

[5] D. Barden, H. Le, and M. Owen. Limiting behaviour of Fréchet means in the space of phylogenetic trees. *arXiv preprint arXiv:1409.7602*, 2014.

[6] M. Berger. *A panoramic view of Riemannian geometry*. Springer Science & Business Media, 2003.

[7] L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics* 27(4):733–767, 2001.

[8] A. Borbély. Some results on the convex hull of finitely many convex sets. *Proceedings of the American Mathematical Society* 126(5):1515–1525, 1998.

[9] B. H. Bowditch. Some results on the geometry of convex hulls in manifolds of pinched negative curvature. *Commentarii Mathematici Helvetici* 69(1):49–81, 1994.

[10] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, vol. 319. Springer Science & Business Media, 1999.

[11] J. Canny and J. Reif. New lower bound techniques for robot motion planning problems. *28th Annual Symp. on Foundations of Computer Science(FOCS)*, pp. 49–60, 1987.

[12] T. M. Chan. Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discrete & Computational Geometry* 16(4):361–368, 1996.

[13] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete & Computational Geometry* 6:485–524, 1991, doi:10.1007/BF02574703.

[14] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete and Computational Geometry* 10(1):377–409, 1993.

[15] J. Chen and Y. Han. Shortest paths on a polyhedron, Part I: Computing shortest paths. *International Journal of Computational Geometry & Applications* 6(02):127–144, 1996, doi:10.1142/S0218195996000095.

[16] V. Chepoi. Graphs of some CAT(0) complexes. *Advances in Applied Mathematics* 24(2):125–179, 2000.

[17] V. Chepoi, F. F. Dragan, and Y. Vaxès. Distance and routing labeling schemes for non-positively curved plane graphs. *Journal of Algorithms* 61(2):60–88, 2006.

[18] V. Chepoi and D. Maftuleac. Shortest path problem in rectangular complexes of global nonpositive curvature. *Computational Geometry* 46(1):51–64, 2013, doi:10.1016/j.comgeo.2012.04.002.

[19] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1963.

[20] M. Dror, A. Efrat, A. Lubiw, and J. S. Mitchell. Touring a sequence of polygons. *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 473–482, 2003, doi:10.1145/780542.780612.

[21] M. Elder and J. McCammond. CAT(0) is an algorithmic property. *Geometriae Dedicata* 107(1):25–46, 2004, doi:10.1023/B:GEOM.0000049096.63639.e3.

[22] J. Felsenstein. Confidence limits on phylogenies: an approach using the bootstrap. *Evolution* pp. 783–791, 1985.

[23] P. Fletcher, J. Moeller, J. Phillips, and S. Venkatasubramanian. Horoball hulls and extents in positive definite space. *Algorithms and Data Structures*, pp. 386-398. Springer Berlin Heidelberg, Lecture Notes in Computer Science 6844, 2011, doi:10.1007/978-3-642-22300-6_33.

[24] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Advances in applied mathematics* 38(3):302–323, 2007.

[25] M. Gromov. Hyperbolic groups. *Essays in Group Theory*, pp. 75-263. Springer New York, Mathematical Sciences Research Institute Publications 8, 1987, doi:10.1007/978-1-4613-9586-7_3.

[26] L. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica* 2(1-4):209–233, 1987, doi:10.1007/BF01840360.

[27] L. J. Guibas and J. Hershberger. Optimal shortest path queries in a simple polygon. *Journal of Computer and System Sciences* 39(2):126 – 152, 1989, doi:10.1016/0022-0000(89)90041-X.

[28] F. Haglund and D. T. Wise. Special cube complexes. *Geometric and Functional Analysis* 17(5):1551–1620, 2008.

[29] J. Hershberger and J. Snoeyink. Computing minimum length paths of a given homotopy class. *Computational Geometry* 4(2):63–97, 1994, doi:10.1016/0925-7721(94)90010-8.

[30] J. Hershberger and S. Suri. An optimal algorithm for euclidean shortest paths in the plane. *SIAM J. Comput* 28:2215–2256, 1997.

[31] S. Holmes. Statistical approach to tests involving phylogenies. *Mathematics of evolution and phylogeny. Oxford University Press, Oxford, UK* pp. 91–120, 2005.

[32] M. Ishaque and C. D. Tóth. Relative convex hulls in semi-dynamic arrangements. *Algorithmica* 68(2):448–482, 2014, doi:10.1007/s00453-012-9679-6.

[33] N. Karmarkar. A new polynomial-time algorithm for linear programming. *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pp. 302–311, 1984.

[34] L. G. Khachiyan. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics* 20(1):53–72, 1980.

[35] D. G. Kirkpatrick and R. Seidel. The ultimate planar convex hull algorithm? *SIAM Journal on Computing* 15(1):287–299, 1986.

[36] D. Lee and F. P. Preparata. Euclidean shortest paths in the presence of rectilinear barriers. *Networks* 14(3):393–410, 1984, doi:10.1002/net.3230140304.

[37] D.-T. Lee and B. J. Schachter. Two algorithms for constructing a Delaunay triangulation. *International Journal of Computer and Information Sciences* 9(3):219–242, 1980.

[38] B. Lin, B. Sturmfels, X. Tang, and R. Yoshida. Convexity in tree spaces. *arXiv preprint arXiv:1510.08797*, 2015.

[39] D. Maftuleac. Algorithms for distance problems in planar complexes of global nonpositive curvature. *International Journal of Computational Geometry & Applications* 24(01):1–38, 2014, doi:10.1142/S0218195914500010.

[40] E. Miller, M. Owen, and J. S. Provan. Polyhedral computational geometry for averaging metric phylogenetic trees. *Advances in Applied Mathematics* 68:51–91, 2015.

[41] J. S. Mitchell. Geometric shortest paths and network optimization. *Handbook of Computational Geometry*, pp. 633–701. Elsevier Science Publishers B.V. North-Holland, 1998.

[42] J. S. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing* 16(4):647–668, 1987.

[43] D. M. Mount. Storing the subdivision of a polyhedral surface. *Discrete & Computational Geometry* 2(1):153–174, 1987, doi:10.1007/BF02187877.

[44] T. M. Nye. An algorithm for constructing principal geodesics in phylogenetic treespace. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on* 11(2):304–315, 2014.

[45] M. Owen. Computing geodesic distances in tree space. *SIAM Journal on Discrete Mathematics* 25(4):1506–1529, 2011.

[46] M. Owen and J. S. Provan. A fast algorithm for computing geodesic distances in tree space. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)* 8(1):2–13, 2011.

[47] F. Ronquist and J. P. Huelsenbeck. MrBayes 3: Bayesian phylogenetic inference under mixed models. *Bioinformatics* 19(12):1572–1574, 2003.

[48] R. Seidel. Convex hull computations. *Handbook of Discrete and Computational Geometry (2nd edition)*, pp. 495–512. CRC Press, Inc., 2004.

[49] G. Toussaint. An optimal algorithm for computing the relative convex hull of a set of points in a polygon. *Signal Processing III: Theories and Applications, Proc. EURASIP-86, Part 2*, pp. 853–856, 1986.

[50] G. Toussaint. Computing geodesic properties inside a simple polygon. *Rev. Intell. Artif.* 3:9–42, 1989.

[51] J. W. Tukey. Mathematics and the picturing of data. *Proceedings of the International Congress of Mathematicians*, vol. 2, pp. 523–531, 1975.

[52] G. U. Yule. A mathematical theory of evolution, based on the conclusions of Dr. JC Willis, FRS. *Philosophical Transactions of the Royal Society of London. Series B, Containing Papers of a Biological Character* pp. 21–87, 1925.